

**Technical Report
TR-1104**

Activity Level Change Detection for Persistent Surveillance

**F. Liu
L.A. Bush**

**2 October 2004
Issued 26 May 2006**

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS



Prepared for the Department of the Air Force under Contract FA8721-05-C-0002.

Approved for public release; distribution is unlimited.

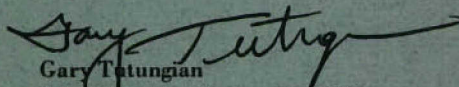
This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This work was sponsored by the Air Force Space and Missile Systems Center, YS, under Contract FA8721-05-C-0002.

This report may be reproduced to satisfy needs of U.S. Government agencies.

The ESC Public Affairs Office has reviewed this report, and it is releasable to the National Technical Information Service, where it will be available to the general public, including foreign nationals.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER


Gary Tutungian
Administrative Contracting Officer
Plans and Programs Directorate
Contracted Support Management

Non-Lincoln Recipients

PLEASE DO NOT RETURN

Permission has been granted to destroy this document, when it is no longer needed.

Massachusetts Institute of Technology
Lincoln Laboratory

Activity Level Change Detection for Persistent Surveillance

*F. Liu
Kirajalein
L.A. Bush
Group 104*

Technical Report 1104

2 October 2001
Issued 26 May 2006

Approved for public release; distribution is unlimited.

ABSTRACT

A new approach to GMTI data exploitation for large area persistent surveillance is presented. Instead of traditional target tracking, this approach utilizes GMTI data as moving spots on the ground to estimate the level of activities and detect unusual activities such as military deployments.

A multilayer hierarchical exploitation scheme is proposed. This computational framework has clean interfaces between layers consisting of multiple processing modules. Various data processing, machine learning, and reasoning algorithms can be implemented in these modules. This system is easily extendable and can be tested using a generalized test bed.

The development of two processing modules, vehicular volume and convoy detector, is described. For the vehicular volume module, US highway data were used as a surrogate of long-term GMTI surveillance data. The relationship between the activity level of Norfolk Naval Base and the traffic pattern on a road leading to the Base is studied. The convoy detection module, developed using real GMTI data, contains an algorithm that detects convoys without explicit target tracking.

An end-to-end testing facility was also developed. Using this test bed, the system can be tested at different levels: as an individual processing module, as multiple cooperating processing modules across layers, or as the entire system.

ACKNOWLEDGEMENTS

The authors would like to thank Catherine McGhee and Prof. Brian Smith for the access to the ADMS Virginia traffic database, Harold Heggstad for obtaining the Norfolk Naval Base ship activity historical data, Peter Mastromarino for the dissimilarity computation in Section 3.9, and Jason Oda for many of the graphics.

TABLE OF CONTENTS

Abstract	iii
Acknowledgments	v
List of Illustrations	ix
List of Tables	xi
 1. INTRODUCTION	 1
2. A FRAMEWORK FOR MULTISENSOR PERSISTENT SURVEILLANCE	3
3. NORFOLK TRAFFIC STUDY	7
3.1 Norfolk Traffic Data	7
3.2 Processing Module Flowchart	8
3.3 Feature Extraction	8
3.4 Feature Refinement	10
3.5 Feature Modeling	12
3.6 Feature Evaluation	14
3.7 Impact of Ship Sizes	15
3.8 Analysis of I-564 EB Vehicular Volume Data	20
3.9 Model Update	22
 4. CONVOY DETECTION USING GMTI DATA	 27
4.1 Identify Convoy Candidates in a Scan	28
4.2 Motion Model	29
4.3 Convoy Candidate Correlation Between Two GMTI Scans	30
4.4 Evidence Evaluation	31
 5. A TEST BED FOR EVALUATING DETECTION AND REASONING ALGORITHMS	 33
5.1 Vehicle Sources	34
5.2 Sensors	34
5.3 Detector	35
5.4 User Interface	35
5.5 Software Architecture	36
 6. SUMMARY	 39

LIST OF ILLUSTRATIONS

Figure No.		Page
1	Taiwan scenario force mobilization time line.	4
2	Taiwan scenario situational analysis.	4
3	Computational framework for persistent surveillance using multisensor data, illustrated using Taiwan scenario.	5
4	Functional structure of a processing module.	6
5	Hampton Roads area. I-564 leads to Norfolk Naval Base.	7
6	Traffic data processing module flowchart.	8
7	Means and standard deviations of vehicular volume. Shown are mean (solid blue) and standard deviation (shaded blue) of 31 normal days as well as mean (solid red) and standard deviation (shaded red) of 10 abnormal days.	9
8	Means and standard deviations of vehicular speed. Shown are mean (solid blue) and standard deviation (shaded blue) of 31 normal days as well as mean (solid red) and standard deviation (shaded red) of 10 abnormal days.	10
9	Student's <i>t</i> -test between normal and abnormal day vehicular volume data at level 0.05.	11
10	Student's <i>t</i> -test between normal and abnormal day vehicular speed data at level 0.05.	11
11	Time interval with the best separation between normal and abnormal day vehicular volume data.	12
12	Vehicular volume is assumed to have a Gaussian distribution.	13
13	Histograms of normal day vehicular volume data between 0900 and 1100 hours. Each picture contains six curves. Each curve corresponds to a five-minute mark.	13
14	Results of vehicular volume feature evaluation. Each ROC curve is generated using 30 minutes aggregated test results.	14
15	ROC curve generated from 2-hour aggregated results. To detect 90% of all anomalies, 9 out of 20 detections would be false alarms.	15

LIST OF ILLUSTRATIONS (CONTINUED)

Figure No.		Page
16	Vehicular volume means and standard deviations of normal, small ship, and large ship days at 1100 hours.	17
17	Vehicular volume normal day means and standard deviations of small and large ship days.	17
18	Vehicular volume means (top) and Student's <i>t</i> -test results (bottom) for 31 normal and 6 small ship days.	18
19	Vehicular volume 2-hour aggregation ROC for small ship days.	19
20	Vehicular volume means (top) and Student's <i>t</i> -test results (bottom) for 31 normal and 8 large ship days.	19
21	Vehicular volume 2-hour aggregation ROC for large ship days.	20
22	I-564 EB vehicular volume means (top) and Student's <i>t</i> -test results (bottom) for 31 normal and 3 abnormal days.	21
23	I-564 EB vehicular volume 1.5-hour aggregation ROC.	21
24	Vehicular volume similarity between the same months of adjacent years.	23
25	Vehicular volume similarity between adjacent months.	23
26	Averaged vehicular volume similarity between two months separated from 1 to 36 months.	24
27	Monthly vehicular volume similarity comparison using a chi-square dissimilarity measure. The two months compared are separated from 1 to 36 months.	25
28	Flowchart of convoy detection algorithm for GMTI data.	28
29	Identify convoy candidates by clustering GMTI data. The gray lines form a road intersection. 1) Raw GMTI detections. 2) Isolated detections are removed. 3) Six clusters are found initially. 4) One qualified convoy candidate cluster remains after refinement. The shape of the clusters, the distance between the clusters, and the consistency of the detections in each cluster are examined in the process.	29

LIST OF ILLUSTRATIONS (CONTINUED)

Figure No.		Page
30	Convoy motion model. The apparent cluster speed V_R along the radar range direction is back-projected to the estimated cluster heading direction (blue line) to estimate the cluster speed V .	30
31	Correlating convoy candidate clusters. The candidate cluster from the new GMTI scan is shown in brilliant colors, and the one carried over from the previous scan is in faded color. The position of the candidate from the previous scan is predicted in the new scan. Euclidean distance between the “predicted position” of the old candidate and the “new position” of the new candidate is used to correlate the candidates.	31
32	Top-level diagram for real-time anomaly detection in simulated traffic.	34
33	The structure and the learning scheme of simplified fuzzy ARTMAP classifier.	35
34	Test bed graphical user interface. The three bottom panels control the red, white, and blue vehicle sources. The upper-right panel displays the traffic flow and sensor locations (tripwires shown). The upper-left panel shows the internal parameters such as the feature vector values and the status of the traffic (normal or abnormal). The middle-left panel is the control for SFAM.	36
35	Test bed “publish-subscribe” software design. The upper-left part hosts the sensors. The upper-right controls the tabular displays of parameters. The lower-left contains the detectors. The lower-right provides all the graphical displays.	37

LIST OF TABLES

Table No.		Page
1	Truth Data for Small and Large Ship Arrival and Departure: Six Small Ship Days and Eight Large Ship Days	16

1. INTRODUCTION

The Space-Based Radar (SBR) system aims to provide persistent surveillance at a global scale. The two principal operational modes of the system are the Synthetic Aperture Radar (SAR) and the Ground Moving Target Indicator (GMTI). While SAR provides information on nonmoving ground scatterers, GMTI data contains information on moving targets. Due to the large flow of data that will be generated by the SBR sensors, the automatic exploitation of the sensor data is a key component of the SBR system.

This report presents work in progress in the area of SBR GMTI data exploitation. Since tracking at a global scale is prohibitively expensive for SBR GMTI sensors, we explore methods that can extract beneficial information from the uncorrelated GMTI data. One approach is to model the behavior of the GMTI “spots” under normal circumstances to detect unusual activities such as military deployments. By this approach, we view the SBR GMTI exploitation as a pattern recognition, machine learning, and data mining problem.

We propose here a multilayer hierarchical exploitation scheme that has clean interfaces between layers and can be easily extended. Each layer contains multiple processing modules, and all of the modules in the system have a similar functional structure in terms of data analysis, pattern modeling, and anomaly detection. This system can be viewed as a computational framework for multisensor persistent surveillance.

This report describes the development of two processing modules for the proposed system: vehicular volume and convoy detector. Within the framework of activity level anomaly detection, both modules generate low-level computational features. Since GMTI surveillance data over a prolonged period of time is not readily available, traffic data collected in Virginia Hampton Roads area are used as a surrogate data source for the development of the vehicular volume feature. The relationship between the activity level of Norfolk Naval Base and the traffic pattern on a road leading to the Base is studied in detail. Convoys are important indicators of military movements. To capture this information, an algorithm for detecting convoys in GMTI data without explicit target tracking is also developed.

An end-to-end testing facility has also been developed in this work. The proposed computational framework has the ability to incorporate many processing modules that are specialized in exploiting different aspects of sensor data. These modules can be implemented using a variety of data processing and pattern analysis techniques. When building such a system, a test environment for testing and evaluating the processing modules is highly desirable. This flexible test bed is developed to facilitate the system development at various levels, should it be an individual module or a set of cooperating modules.

In the following section, the proposed framework for multisensor persistent surveillance is presented. In Section 3, the Norfolk traffic study is used as an example to explain the inner working of the framework. Section 4 explains the convoy detection algorithm. The test bed for system testing and evaluation is presented in Section 5.

2. A FRAMEWORK FOR MULTISENSOR PERSISTENT SURVEILLANCE

The data processing and reasoning scheme developed in this work for activity change detection can be considered generally as a framework for multisensor large-area persistent surveillance. This framework has a multilevel bottom-up hierarchy, where the sensor data are at the lowest level. The entire system consists of highly modular components that can be extended easily to accommodate new sensor inputs and new processing and reasoning schemes.

To help explain this system design, we first take a look at how human analysts would conduct a situational analysis on an example surveillance scenario. The setting of the scenario is the Taiwan Strait. The question to address is whether or not Mainland China is preparing an imminent attack on Taiwan. To answer this question from a surveillance point of view, we consider how the military activity level in that area would change if the preparation for the attack is in motion.

Figure 1 shows a possible military preparation timeline starting from six months prior to the final attack. Potential changes in site activities are listed by military branches. We are particularly interested in the changes that occur three to six months ahead of the attack. Around that timeframe, some early site-level activity changes may be present. For instance, increased ship and pier retrofit activities at navy shipyards and ports; increased activity level at army garrisons, live-fire ground weapon ranges, and marshaling yards; as well as increased activities at Air Force main operating bases (MOBs) and live-fire air weapon ranges.

Let us use the Army marshaling yards as an example to continue the analysis. To decide if the activity level in a marshaling yard is unusual, the vehicular volume, number of train cars, train speed, number of convoys, etc., can be monitored. Since the information on these items can be obtained readily from the sensor data, they are regarded as low-level observables. These observables form the foundation of our analysis.

Figure 2 structures the analysis of the Taiwan scenario into a multilevel situational analysis flowchart. To answer the ultimate question at the top level, we monitor the activities at various mid-level sites. To determine the level of activity at a particular site, we keep tracking relevant low-level observables. In essence, this flowchart is built from the top down by a knowledge-driven analysis process. The result is a multilevel hierarchical system. Using this system to determine if China is preparing an attack on Taiwan, the analysts would first gather information from the low-level observables to reach conclusions at particular sites and then use the results from various sites to assess the situation of the entire region.

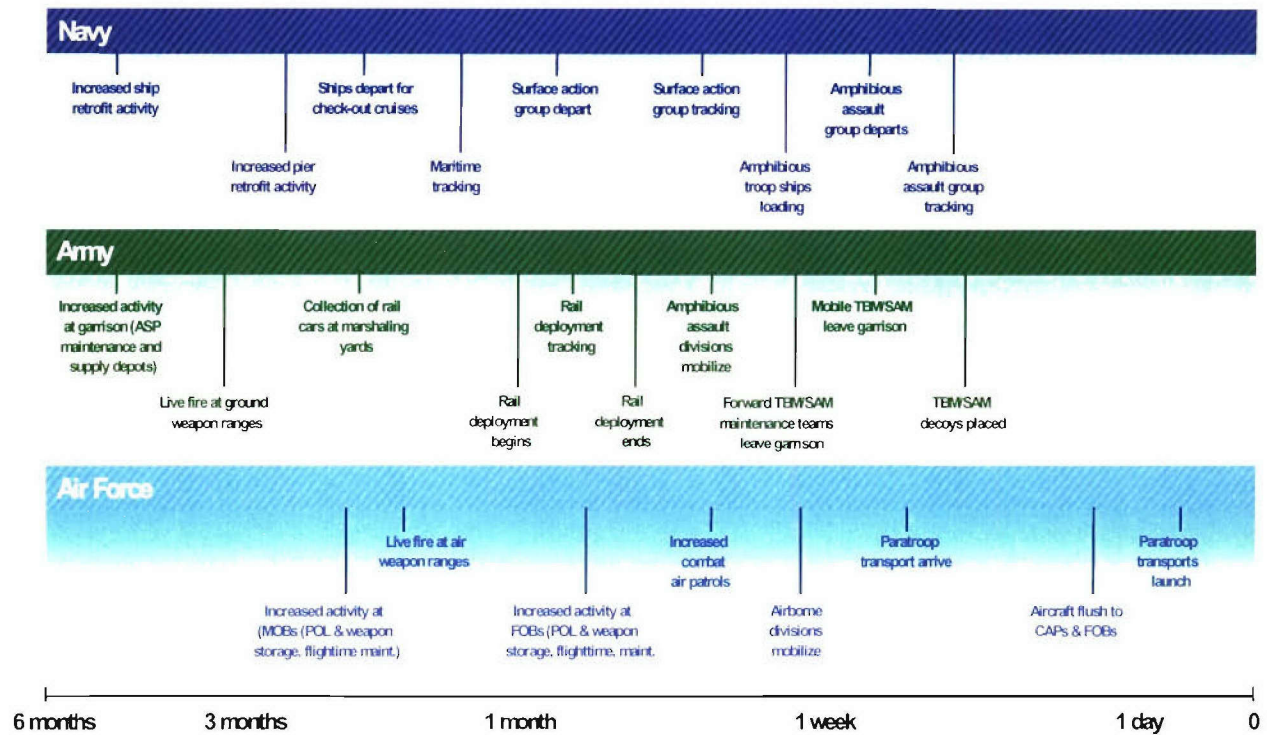


Figure 1 Taiwan scenario force mobilization time line.

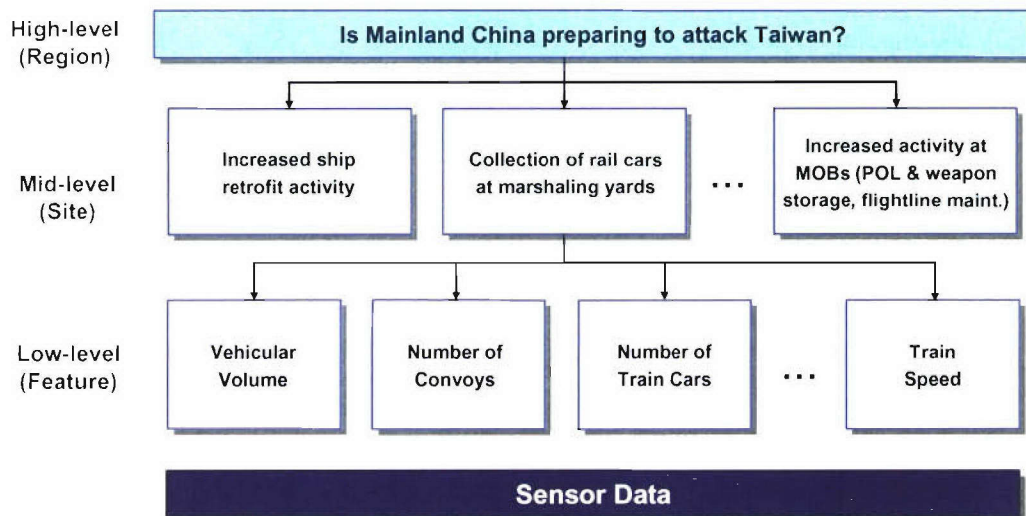


Figure 2. Taiwan scenario situational analysis.

As mentioned previously, sensor data can provide information on low-level observables. To build a computational system that uses sensor data as inputs, we reverse the analysis process in Figure 2. The result is a data-driven bottom-up hierarchy, shown in Figure 3. At the entry level, features are extracted from sensor data to support the low-level inference. The results are then sent to the next level up to ascertain the state of site activities. Finally, the results from all relevant sites are used to reach a regional conclusion.

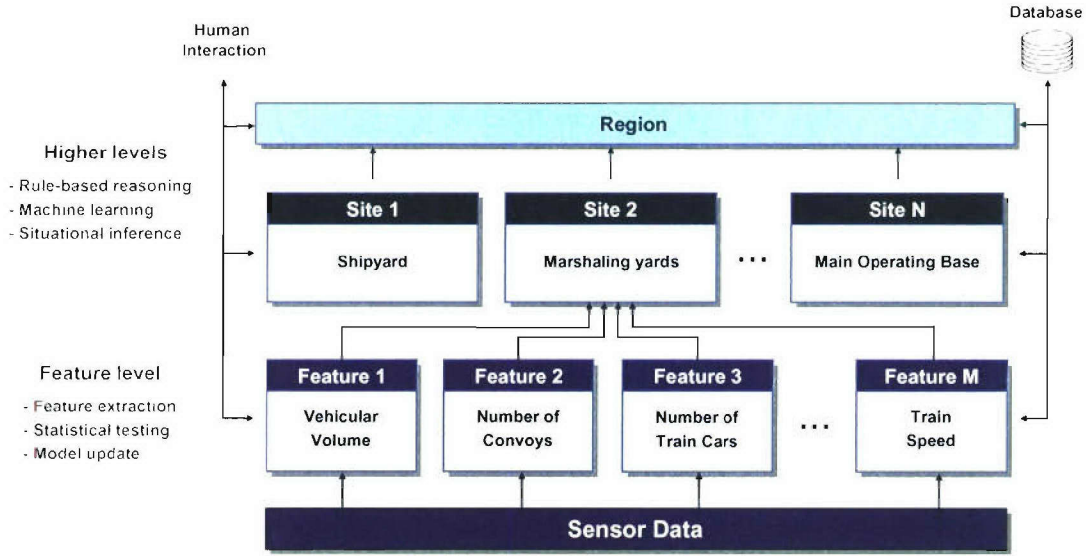


Figure 3. Computational framework for persistent surveillance using multisensor data, illustrated using Taiwan scenario.

The basic elements of this computational system are its processing modules. At the core of these modules are statistical models and learning schemes that are used to evaluate incoming data for their normalcy. Different algorithms can be implemented in the modules. A simple module, typically a low-level one, can contain a straightforward statistical model such as the Gaussian model. A more sophisticated module can be a rule-based expert system, a pattern recognition algorithm, a machine-learning scheme, or a hybrid system.

Although the computational models implemented in the processing modules may take many different forms, the design of the modules can share a common functional structure. Illustrated in Figure 4, a module first extracts appropriate computational features from the input data. Whenever necessary, the features are refined. The feature values are used to build or update the computational model in the module. Finally, the feature values are evaluated against the existing model to determine whether they are normal or not. The common structure of the modules facilitates system testing and extension.




Figure 4. Functional structure of a processing module.

This multilevel hierarchical computational framework emphasizes evidence accumulation and continuous learning. Valuable information, including model parameters, intermediate results, and previous conclusions and observations, are stored in the database. At every level of reasoning, human expertise and intervention are an integrated part of the system. It can be easily extended to accept new inputs, generate new features, monitor additional sites, and ultimately provide more surveillance coverage.

In the next two sections, we use highway traffic volume and convoy detection as examples to explain how entry-level feature modules can be built and used to provide information for higher-level inference.

3. NORFOLK TRAFFIC STUDY

To study the activity patterns in a region, we need a data set that spans over a long period of time. Since a suitable GMTI data set is not available to us, we used US highway traffic data as a surrogate. The goal is to explore how certain attributes of the traffic data can be used to monitor the activity level in a geographic area. In this section, the construction of a processing module designed to utilize the vehicular volume as a low-level feature is explained step by step.

3.1 NORFOLK TRAFFIC DATA

Data used in this study were downloaded from the Archived Data Management System for Virginia (ADMS Virginia). This traffic database is sponsored by the Federal Highway Administration and Virginia Department of Transportation. It is currently managed by the Smart Travel Lab of the University of Virginia. Traffic data are collected using embedded magnetic loop sensors located throughout the Virginia Hampton Roads area, shown in the left-hand image of Figure 5. Vehicular speed, volume, and occupancy data are collected every 20 seconds and then aggregated and recorded every minute. Due to high failure rates of the sensors, data screening is essential.

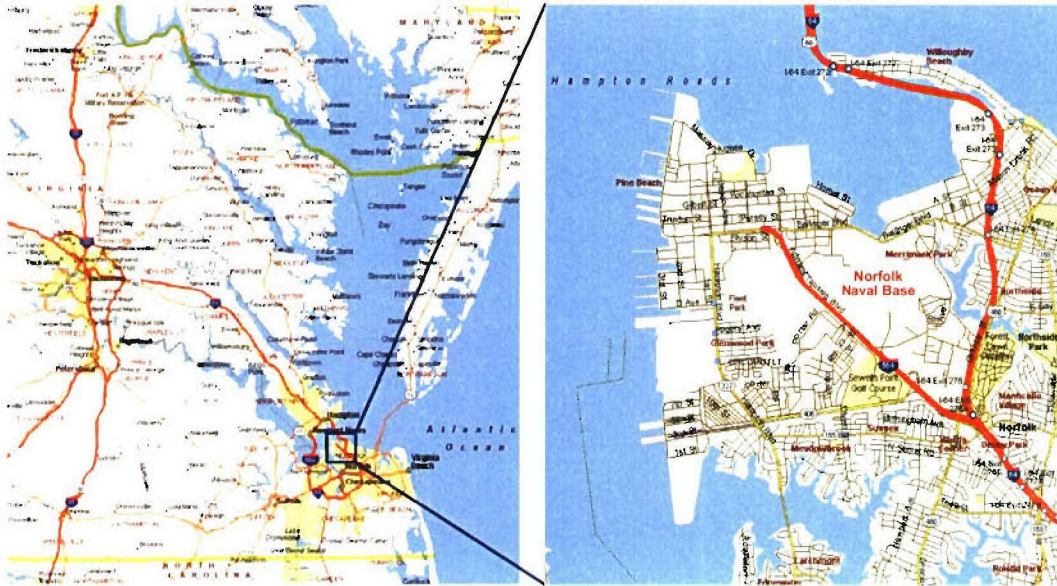


Figure 5. Hampton Roads area. I-564 leads to Norfolk Naval Base.

In addition to the traffic data, ground truth of regional activity level is necessary for algorithm development and testing. The truth data, however, is not available in the ADMS Virginia database. This issue was resolved by inferring the truth from the ship arrival and departure activities at Norfolk Naval Base. The right-hand image of Figure 5 shows the road details of the Norfolk area. The major highway that leads to the Naval Base is Route I-564. As expected, major Base activities tend to have a significant impact on the traffic pattern of the nearby highways, including I-564.

In this study, the regular weekday traffic is considered as the “normal” activity, and the traffic during the ship arrival and departure days are regarded as “anomalies.” Initially, data collected at Station 131 on I-564 westbound (WB) was used as the main data source. The vehicular volume at this station was studied extensively and used as a feature-level indicator to Base activities. Later, data collected at I-564 eastbound (EB) Station 135 are processed to study traffic patterns across months, seasons, and years.

At the beginning of the study, the truth data was obtained from the newspaper webpage <http://www.hamptonroads.com/military/homecomings>. (This truth data was later confirmed by a ship arrival and departure list obtained from the Norfolk Naval Base.) The available truth data limited the time span of the traffic data used in the study to about three months. We screened the weekday vehicular volume and speed data collected between April 7 and July 3, 2003 at Station 131 and found 41 days of usable data. Among the 41 days, ten are considered abnormal. This 41-day data set constitutes the initial Norfolk traffic data set.

3.2 PROCESSING MODULE FLOWCHART

The computational flowchart for the traffic volume analysis module is shown in Figure 6, where the “Feature Refinement” step is specified as a gating function. Each step of the processing is explained in detail in the following subsections.

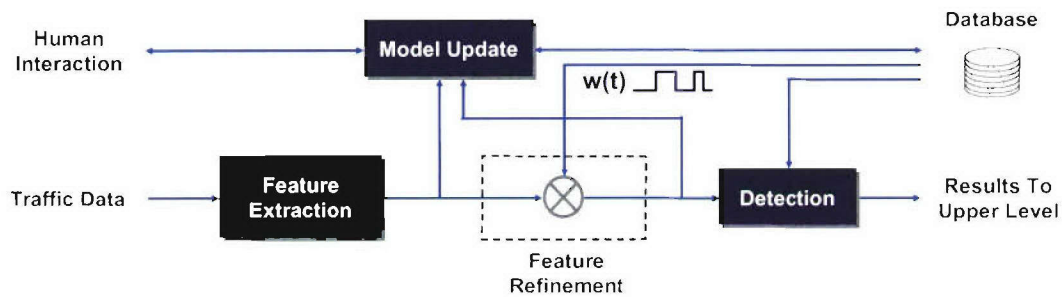


Figure 6. Traffic data processing module flowchart.

3.3 FEATURE EXTRACTION

The typical first step of a module is to preprocess the incoming data and extract features that are relevant to the subsequent processing. For Norfolk traffic data analysis, the data screening mentioned

previously can be considered as preprocessing. The screened per-minute data typically contain some high frequency fluctuations that can be regarded as noise. To reduce the fluctuation, we filtered data at every five-minute mark using a moving average window that averages the data over the past ten minutes. For each 24-hour period starting at 00:10 hour, this process results in 287 data points, which are used as computational features in the subsequent processing.

Can the vehicular volume and speed features indicate the level of Base activities? To answer the question, the means and standard deviations of these features are first computed at each five-minute mark over the 31 normal and 10 abnormal days, respectively. The results are shown in Figures 7 and 8. The solid lines in the figures represent the mean values, and the boundaries of the shaded areas correspond to the standard deviations. Since I-564 WB is inbound to Norfolk Naval Base, the traffic volume peaks in the morning as expected.

To be able to distinguish the abnormal traffic pattern from the normal ones, we would like to see good separations between the normal and abnormal curves in the two figures. Taking into account the variances of the data, the best separation appears to be in the vehicular volume, during the morning period. This tells us that, to use vehicular volume and speed as computational features for anomaly detection, we need to refine the features to specify the time intervals in a day when the features are the most discriminative. This refinement is represented by the gating function $w(t)$ in Figure 6.




Figure 7 Means and standard deviations of vehicular volume. Shown are mean (solid blue) and standard deviation (shaded blue) of 31 normal days as well as mean (solid red) and standard deviation (shaded red) of 10 abnormal days.




Figure 8. Means and standard deviations of vehicular speed. Shown are mean (solid blue) and standard deviation (shaded blue) of 31 normal days as well as mean (solid red) and standard deviation (shaded red) of 10 abnormal days.

3.4 FEATURE REFINEMENT

For the vehicular volume and speed data, the separations between the normal and the abnormal data sets are quantified by using the Student's t -test. The Student's t -test determines whether two sets of samples are drawn from the same statistical distribution. Therefore, it is a test between the two hypotheses: H_0 : two sample sets are from the same distribution; and H_1 : two sample sets are from different distributions. The sample sets used in our tests are populated by the vehicular volume or speed data from the "normal" and the "abnormal" days. Figures 9 and 10 show the results of the Student's t -tests. The "significance" values in the figures are the probabilities that the two sample sets are from the same distribution. The tests were conducted at level 0.05, which means that a 5% threshold for the significance values was used to choose between the two hypotheses. In the two figures, H_0 and H_1 are marked in red with value 0 and 1, respectively.

The Student's t -test results show that the normal day and the abnormal day vehicular volume data have the best separations in the early part of a day, between 0830 and 1230 hours. In Figure 11, this time interval is superimposed onto the vehicular volume mean and standard deviation data shown in Figure 7. The Student's t -tests also show that the speed data do not have consistent separations. Therefore, this particular vehicular speed data set is not a strong indicator for the Norfolk Naval Base activities. In the following analysis, we will focus on only the volume data.




Figure 9. Student's t-test between normal and abnormal day vehicular volume data at level 0.05.




Figure 10. Student's t-test between normal and abnormal day vehicular speed data at level 0.05.




Figure 11. Time interval with the best separation between normal and abnormal day vehicular volume data.

3.5 FEATURE MODELING

Computational models are at the core of all processing modules. These models capture the behavior of the input features under normal circumstances. During operation, the incoming feature values are evaluated against the models to detect an anomaly.

For the vehicular volume feature, we assumed Gaussian models for data points collected at the same five-minute mark across all “normal” days. This assumption is illustrated in Figure 12. One way to verify the Gaussian assumption is to visually examine the histograms of the normal day vehicular volume data. Based on the results of the feature refinement, we narrowed our focus to the time interval between 0900 and 1100 hours for modeling. Histograms of vehicular volume are generated at every five-minute mark for the two-hour period. The results are plotted in the four pictures of Figure 13. Each picture contains six separate curves. These histograms indicate that Gaussian distribution is an appropriate assumption for the volume data. The Gaussian model parameters are estimated from the data as the statistical sample means and standard deviations.




Figure 12. Vehicular volume is assumed to have a Gaussian distribution.

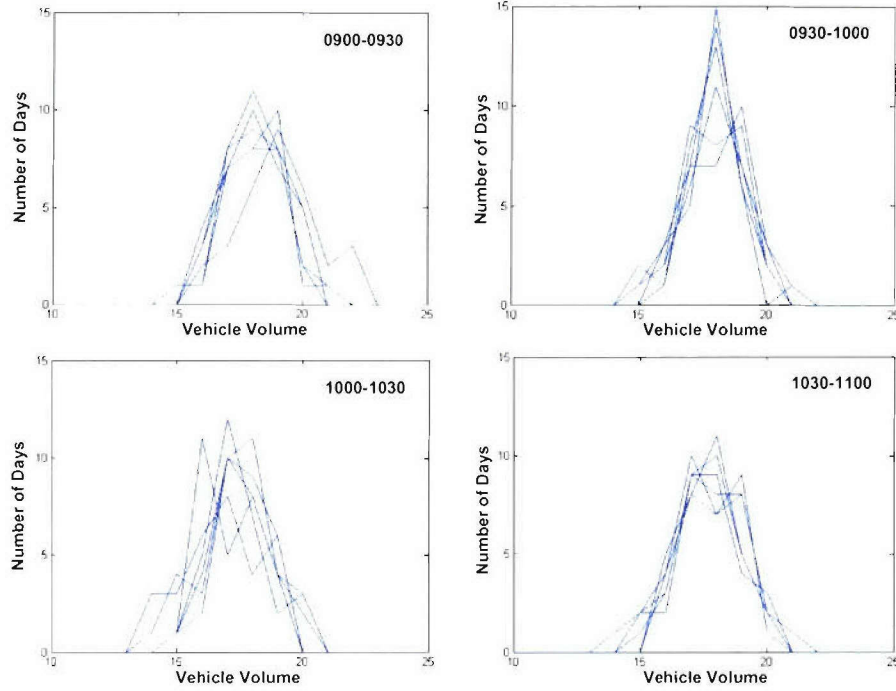


Figure 13. Histograms of normal day vehicular volume data between 0900 and 1100 hours. Each picture contains six curves. Each curve corresponds to a five-minute mark.

3.6 FEATURE EVALUATION

The last processing step in a module is to evaluate the incoming feature values against the models established for the “normal” activity patterns. The results help to determine if an anomalous situation is present.

For the vehicular volume feature, the evaluation is conducted at every five-minute mark between 0900 and 1100 hours. The models for the volume data are one-dimensional Gaussians and the anomalies are expected to increase the traffic volume. Therefore, a single-valued decision boundary is used to decide if a volume data value is within the “normal” range. At each five-minute mark, all 41 available data values (31 from normal days and 10 from abnormal days) are used in testing. The probabilities of detection (P_D) and false alarm (P_F) are calculated by comparing the test outcome against the truth data. By varying the value of the decision boundary, a set of P_D 's and P_F pairs are generated and then plotted as Receiver Operating Characteristics (ROC) curves. In this work, the probabilities are first calculated by aggregating the test results in four 30-minute intervals. The resulting ROC curves are shown in Figure 14. Then the results from the entire two-hour time interval are aggregated to generate the ROC curve in Figure 15.




Figure 14. Results of vehicular volume feature evaluation. Each ROC curve is generated using 30 minutes aggregated test results.




Figure 15. ROC curve generated from 2-hour aggregated results. To detect 90% of all anomalies, 9 out of 20 detections would be false alarms.

Note that, in this experiment, the same set of “normal” day data was used for both model generation and testing. A more sophisticated test method, such as “leave-one-out,” can be used instead. Nevertheless, since the statistical distribution of the normal day data is close to Gaussian, a significant difference in the test results due to the Gaussian assumption is unlikely.

An effective and efficient system should operate at locations close to the upper left corner of the ROC plot, i.e., operating with high probability of detection and low probability of false alarm. As shown by the red lines in Figure 15, the performance of the vehicular volume feature is rather poor: to detect 90% of all anomalies, 9 out of 20 detections would be false. In the next subsection, this performance issue is examined in detail.

3.7 IMPACT OF SHIP SIZES

In the process of identifying the main contributors to system performance, we noticed that some of the days had large ships departing or arriving while the others had much smaller ones. Table 1 contains the list of anomalous days provided by the Norfolk Naval Base and the potential numbers of personnel onboard the ships. Presumably, large ships would have more substantial impact on the traffic pattern near the Base than the small ships would.

TABLE 1
Truth Data For Small and Large Ship Arrival and Departure:
Six Small Ship Days and Eight Large Ship Days

	Date	Ship Type and Name	A/D	Crew	Troops
Small Ships	06/11/03	Destroyer Arleigh Burke	Arrival	337	
	06/13/03	Command – Mount Whitney	Arrival	970	
	06/25/03	Amphibious Assault Ship Bataan	Arrival	1,082	
		Amphibious Transport Dock Ponce	Arrival	364	
		Dock Landing Ship Ashland	Arrival	320	
		Dock Landing Ship Gunston Hall	Arrival	320	
	06/26/03	Amphibious Assault Ship Saipan	Arrival	1,067	
	06/27/03	Amphibious Assault Ship Saipan	Arrival	930	
	07/03/03	Crouser Anzio	Arrival	387	
		Destroyer Porter	Arrival	387	
Large Ships	09/22/99	Aircraft Carrier USS Roosevelt	Arrival	6,122	3,061
	02/18/00	Aircraft Carrier USS Dwight Eisenhower	Departure	6,130	3,065
	05/22/00	Aircraft Carrier USS Harry Truman	Departure	6,122	3,061
	06/21/00	Aircraft Carrier USS George Washington	Departure	6,122	3,061
	12/20/02	Aircraft Carrier USS George Washington	Arrival	6,122	3,061
	05/23/02	Aircraft Carrier USS Harry Truman	Arrival	6,122	3,061
	05/29/03	Aircraft Carrier USS Roosevelt	Arrival	6,122	3,061
	05/30/03	Aircraft Carrier USS Roosevelt	Arrival	6,122	3,061

Figure 16 shows the means and standard deviations of traffic volume at 1100 hours during normal, small ship, and large ship days. The overlaps between the normal and the small ship days and between the small and the large ship days are significant. Figure 17 displays the mean of traffic volume in normal days as well as the one-standard-deviation boundaries of the traffic volume in the small and the large ship days. Clearly, the increases of the traffic volume in the small and the large ship days are quite different.




Figure 16. Vehicular volume means and standard deviations of normal, small ship, and large ship days at 1100 hours.




Figure 17. Vehicular volume normal day means and standard deviations of small and large ship days.

To quantify the difference and its effect on system performance, we conducted Student's t -tests, reevaluated the vehicular volume feature, and generated the corresponding ROC plots. The results are shown in Figures 18 through 21.

Figures 18 and 19 show the Student's t -test result and the ROC curve generated using the vehicular volume data from the normal and the small ship days. In this case, the system performance is similar to that in the small and large ship mixed case displayed in Figure 15. In the case of using the normal and the large ship days, shown in Figures 20 and 21, significant improvements are seen for both the shape of the ROC curve and the length of the time interval when the vehicular volume feature remains effective.

By the results shown above, the vehicular volume feature appears to be more effective for detecting the large ship days than the small ship days. This, however, does not mean that the vehicular volume feature is unusable for detecting the small ship days. In fact, an important strength of the proposed computational framework is its potential capability to utilize information provided by multiple imperfect features to produce high quality inference results.




Figure 18. Vehicular volume means (top) and Student's t -test results (bottom) for 31 normal and 6 small ship days.




Figure 19. Vehicular volume 2-hour aggregation ROC for small ship days.




Figure 20. Vehicular volume means (top) and Student's *t*-test results (bottom) for 31 normal and 8 large ship days.




Figure 21 Vehicular volume 2-hour aggregation ROC for large ship days.

3.8 ANALYSIS OF I-564 EB VEHICULAR VOLUME DATA

The vehicular volume data used in the analysis thus far are solely from one collection station on I-564 in the westbound direction. To investigate the consistency of the vehicular volume feature in terms of its correlation to the naval base activity level, data collected at I-564 eastbound Station 135 are also analyzed in this study. This data set consists of 31 normal days and 3 abnormal days. The reduction on the number of abnormal days from the original ten is due to the poor quality of the Station 135 data on the days removed. The analysis procedure used to process the data is the same as the one presented earlier in this section.

The mean vehicular volume of the normal and the abnormal days are plotted in Figure 22. Since I-564 EB is outbound from the naval base, the traffic volume peaks during the afternoon instead of in the morning as the westbound traffic does. Figure 23 shows the system performance ROC curve using results aggregated between 1330 and 1500 hours. This ROC curve is similar to the one shown in Figure 15. Therefore, although the traffic pattern throughout a day is quite different between the eastbound and the westbound data, the vehicular volume feature is consistent in its ability to indicate the activity level at the naval base.

The corresponding vehicular speed data from Station 135 contain mostly failed sensor readings that are not suitable for analysis.




Figure 22. I-564 EB vehicular volume means (top) and Student's *t*-test results (bottom) for 31 normal and 3 abnormal days.




Figure 23. I-564 EB vehicular volume 1.5-hour aggregation ROC

Notice that the peak traffic volume in either direction of I-564 does not present a dramatic change during the ship arrival and departure days. Also, the traffic volume peaks around the same time as the regular rush hours. This is because the release of the military personal from the ships is typically scheduled in the mornings. Since the access road to the Base is gated, it is typically saturated during the peak traffic hours.

3.9 MODEL UPDATE

In Sections 3.5 and 3.6, we discussed how to build a computational model for the vehicular volume feature and how to use the model to evaluate incoming feature values for anomaly detection. Two questions remain: 1) how many different models should a system maintain for this feature; and 2) how frequently should the models be updated? One answer to the first question is that the traffic patterns for weekdays, weekends, and holidays are quite different, and hence separate models are required. A study is now being conducted to answer this question more systematically. In this subsection, we present the analysis results that address the second question.

The data set used in the “model update” analysis consists of weekday vehicular volume data collected from I-564 eastbound Station 135 in years 1999, 2000, 2001, and 2003. (Data from 2002 is incomplete in the ADMS Database.) Known “abnormal” days are excluded. The same screening and smoothing procedures as described in Section 3.3 are used to preprocess the data. For each 24-hour period, this process generates 287 traffic volume values at five-minute marks starting at 00:10 hours.

A series of Student’s *t*-tests are performed to measure the similarity between two monthly vehicular volume data sets. At each five-minute mark, a *t*-test is conducted at 5% significance level. The two test sample sets are compiled respectively from the two monthly data sets. Each sample set contains all the vehicular volume data corresponding to the current five-minute mark. The degree of similarity between the two months is measured as the percentage of the 287 *t*-test results that fail to reject the null hypothesis H_0 . In other words, this similarity measure is the percentage of the 24-hour period when the vehicular volume data from the two samples are statistically similar.

Figure 24 shows the similarity of the vehicular volume pattern between the same months of two adjacent years. Figure 25 shows the similarity between all pairs of adjacent months. By these two figures, it appears that the year-to-year changes in vehicular volume pattern are slightly larger than the month-to-month changes.

To better understand how the monthly vehicular volume pattern changes over the years, Student’s *t*-tests, similar to the ones described above, are performed on pairs of months separated from one to thirty-six months. The degree of similarity is again measured by the percentage of time in a day when the *t*-tests fail to reject the null hypothesis. The results are averaged for each separation interval and then plotted in Figure 26. Apparently, for any particular month, the two most recent months have the most similar traffic volume patterns, followed by the 12th and the 13th months. The similarity also peaks, although less significantly, at the 24th month. The cyclic pattern in the figure reflects seasonal changes in the vehicular volume data.




Figure 24. Vehicular volume similarity between the same months of adjacent years.




Figure 25. Vehicular volume similarity between adjacent months.




Figure 26. Averaged vehicular volume similarity between two months separated from 1 to 36 months.

The comparison of the monthly vehicular volume data is also conducted using a dissimilarity measure. The dissimilarity of the vehicular volume between two months is computed as

$$D = \frac{1}{N} \sum_{i=1}^N \frac{(\bar{x}_i - \bar{y}_i)^2}{\sigma_{x_i}^2 + \sigma_{y_i}^2}$$

where \bar{x}_i and \bar{y}_i are the monthly means at the i^{th} five-minute mark, $\sigma_{x_i}^2$ and $\sigma_{y_i}^2$ are the corresponding variances, and $N = 287$ is the number of five-minute marks in a 24-hour period. When the vehicular volume at each five-minute mark is Gaussian distributed, the difference measure D has a *chi-square* distribution. The results of $(1-D)$ shifted by a constant is averaged for each separation interval and plotted in Figure 27. The scale of the ordinate in the plot is arbitrary.




Figure 27 Monthly vehicular volume similarity comparison using a chi-square dissimilarity measure. The two months compared are separated from 1 to 36 months.

The results in Figures 26 and 27 correspond quite well. Therefore, for the Norfolk traffic data, it is appropriate to update the models for the vehicular volume feature once every other month using the past data. The models for the same month in the past year can also be used as a strong reference.

4. CONVOY DETECTION USING GMTI DATA

In the last section, the functional components of a processing module are explained in detail using the development of the highway vehicular volume feature as an example. As illustrated in Figure 3, the vehicular volume feature is one of the low-level modules in the system. Other modules can be built similarly with the same functional structure shown in Figure 4. For instance, one of the features depicted in Figure 3 is the “Number of Convoys.” Before being counted, however, the convoys need to be distinguished from the rest of the traffic mix. The goal of this section is to construct a module that detects convoys in GMTI data.

This convoy detection algorithm involves first finding groups of detections that are likely to be convoys, then correlating those groups over time, from one GMTI scan to the next. In this manner, the evidence of convoys is accumulated to dismiss false convoys and ultimately to identify the persistent true convoys by applying a threshold to the evidence.

Typically, a GMTI data set consists of information on targets detected in a series of scans over a geographic area. Among other quantities, the data set contains the target position, speed, and the corresponding data quality measures. The scans in a data set, however, may not have been taken close enough in time so that individual vehicles can be tracked easily from one scan to the next. Therefore, the proposed convoy detection algorithm involves first finding groups of detections that are likely to be convoys, then correlating those groups over time, from scan to scan. In this manner, the evidence of a convoy is accumulated to identify the persistent true convoys and dismiss the inconsistent false ones.

The flowchart of this convoy detection algorithm is shown in Figure 28. Convoy candidates are first identified in each scan as qualified clusters of GMTI detections. Each candidate cluster is assigned a **persistence score**, which is increased every time its predicted position in the next new scan has a close match to the position of one of the convoy candidates found in the new scan. The predicted position of a candidate in the next scan is calculated by establishing a motion model for each candidate. When its persistence score is high enough, a candidate is declared a convoy.

Due to the classification of the GMTI data, graphical drawings are used instead of real data in the following to explain the convoy detection algorithm.




Figure 28. Flowchart of convoy detection algorithm for GMTI data.

4.1 IDENTIFY CONVOY CANDIDATES IN A SCAN

The procedure for identifying convoy candidates in GMTI data is essentially a clustering scheme. There are three steps in the process, which are depicted in Figure 29.

The first picture in Figure 29 shows the road segments and the GMTI detections in a fictitious scan. Since isolated detections are unlikely to be part of a convoy, the first step of the processing is to remove them. The result is shown in the second picture. In the next step, clusters are found using a k-means based clustering algorithm. Six such clusters are displayed in the third picture of Figure 29 in different colors. These clusters then undergo several iterations of combining and subdividing so that the resulting clusters are not too close to each other and the members in each cluster are not too far apart. This refinement is necessary since close-by clusters could be a convoy that is split into parts by the clustering algorithm, and far-apart members in a cluster could be unrelated detections. To ensure that a convoy candidate cluster contains multiple detections and has an oblong shape, the number of detections in the cluster, the consistency of the speed value and heading of the detections, and the positions of the detections relative to each other are also examined. The final result of this process is illustrated in the fourth picture.

Once the convoy candidates are identified, each one of them is assigned a persistence score of value 0. In the following, the group motion of each candidate cluster is characterized and used to increase the candidate's persistence score by correlating the candidate with the ones found in subsequent scans.




Figure 29. Identify convoy candidates by clustering GMTI data. The gray lines form a road intersection. 1) Raw GMTI detections. 2) Isolated detections are removed. 3) Six clusters are found initially. 4) One qualified convoy candidate cluster remains after refinement. The shape of the clusters, the distance between the clusters, and the consistency of the detections in each cluster are examined in the process.

4.2 MOTION MODEL

To predict the positions of the candidate clusters in the next GMTI scan, a motion model is generated for each cluster. This process is illustrated in Figure 30.

A convoy candidate cluster contains detections that are closely aligned and have similar speed and heading. In GMTI data, the reported speed for a detected target is the sensor reading of the actual speed along the radar range direction. To predict the position of the cluster in the next scan, the detected cluster speed at the range direction, \mathbf{V}_R , needs to be projected back to the direction of the actual cluster heading. Denote this projected speed as \mathbf{V} . When the quality of the detected speed values in the cluster is reasonably good, \mathbf{V}_R is estimated as the median speed of the detections. Otherwise, it is estimated as the average of the detected speed values in the cluster.




Figure 30. Convoy motion model. The apparent cluster speed V_R along the radar range direction is back-projected to the estimated cluster heading direction (blue line) to estimate the cluster speed V

The actual headings of the GMTI detections are unknown. In this work, the heading of a convoy candidate cluster is estimated by fitting a line to the detections in the cluster using least square regression. The apparent cluster speed V_R along the radar range direction is then back-projected to the fitted line to estimate the actual cluster speed V .

4.3 CONVOY CANDIDATE CORRELATION BETWEEN TWO GMTI SCANS

Once the speed and the heading of convoy candidate clusters in a GMTI scan are estimated, the positions of the clusters in the next scan can be predicted by using these estimates and the elapsed time between the two scans. After a new set of convoy candidates are found in the new scan, the “predicted positions” of the existing candidates are correlated to the “new positions” of the new candidate by calculating their Euclidean distances. This process is shown in Figure 31. The candidate cluster from the new scan is shown in the figure in brilliant colors, and the one carried over from the previous scan is in faded colors.

A strong correlation, i.e., a small Euclidean distance, between a new and an old candidate is considered evidence that the candidates are associated with the same actual convoy. The persistence score of the corresponding candidates increases when such evidence exists.

5. A TEST BED FOR EVALUATING DETECTION AND REASONING ALGORITHMS

In this section, we present a test bed for testing and evaluating the processing modules. As mentioned previously, a variety of statistical modeling, pattern recognition, machine learning, and data mining algorithms can be implemented in the processing modules. During development, various levels of testing are necessary for algorithm modification, parameter adjustment, and performance evaluation. Since the processing modules in this work share a common functional structure, we are able to build a generalized testing system that allows easy change of input, output, and algorithm plug-ins. It also provides access to internal parameter values, which is crucial for testing sophisticated algorithms.

The purpose of this section is to describe the structure of the test bed and to show its capability in performing testing of processing modules. The Vehicular Volume and the Convoy Detection modules developed in Sections 3 and 4 both contain one low-level feature. The system tested in this section is more sophisticated: it uses real-time simulated traffic as input, calculates two sets of low-level features, and then uses these features to learn about the normal traffic pattern at a road intersection by applying a machine learning algorithm. When an unusual traffic pattern emerges, the system flags the operator to report the anomaly in real time.

Figure 32 shows the top-level diagram of the system. Three vehicle sources release vehicles onto a forked section of two-lane roads. Using data collected by the “tripwire” and the “bounding box” sensors, the system computes features for the “detector” to detect anomalies. The detector first learns about the normal traffic patterns from a set of training data. It then examines the incoming data values in real time to detect abnormal patterns. Through a graphical user interface, a human user can manipulate the parameter settings of the vehicle sources, dictate the learning process, and monitor the input, the output, and the internal parameters of the detector.

It should be emphasized that this test bed is designed to be flexible to accommodate a variety of configurations of multilevel hierarchical modules. In our example system, computational features are extracted by the tripwire and the bounding box modules for anomaly detection. When necessary, these features can be replaced or more features can be added. In Figure 32, a “convoy detector” is shown as an example of additional features. To test the system under different input conditions, modules can also be included to transform the source vehicle information. For instance, in the case of SBR, the vehicle source can be transformed into GMTI detections.

The following subsections describe the components of the test bed.

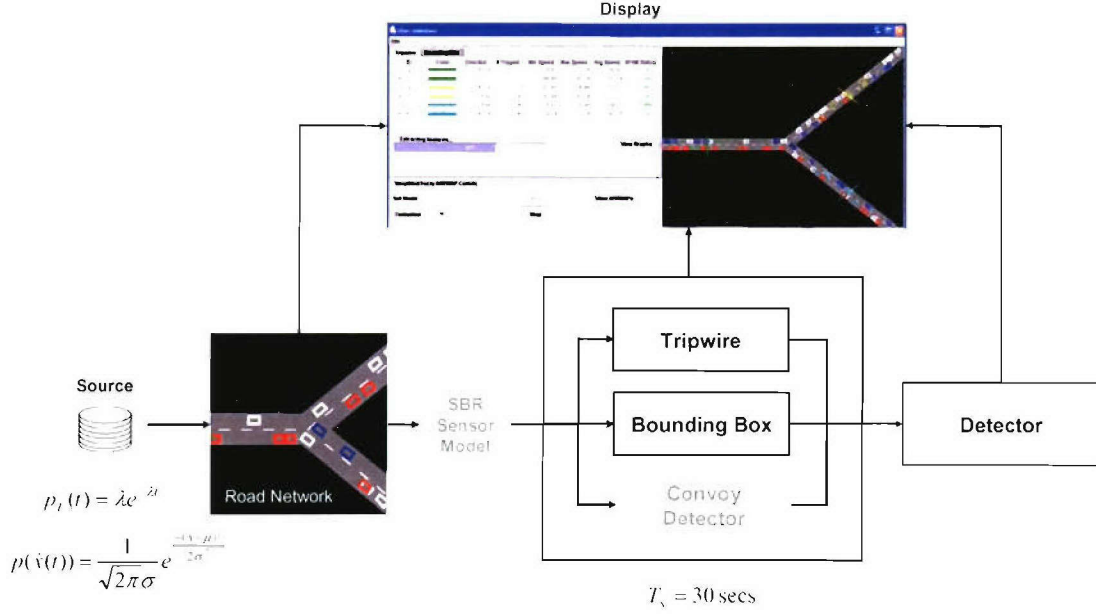


Figure 32. Top-level diagram for real-time anomaly detection in simulated traffic.

5.1 VEHICLE SOURCES

The speeds of the vehicles released by the vehicle sources have Gaussian distributions. The time intervals between two vehicles released at the same source obey Poisson distributions. Faster vehicles can pass the slower ones. Vehicles approaching the intersection choose a direction randomly to proceed, and the choices are uniformly distributed.

5.2 SENSORS

“Tripwire” sensors provide the vehicle heading and volume, as well as the minimum, maximum, and average speed of the vehicles that pass through the tripwire. A “bounding box” encloses an area with traffic activities. It counts the number of vehicles inside the area as well as the ones that cross the borders. Bounding boxes also provide the overall heading and the minimum, maximum, and average speed of the vehicles that are counted. All data are collected every 30 seconds. The different types of data values generated by the sensors are aggregated into vector form and used subsequently as input feature vectors to the detector. The flavor of these feature vectors is consistent with what can be computed from the GMTI data.

5.3 DETECTOR

Data provided by the tripwires and the bounding boxes are used as input to the “detector.” The core of the detector is a neural network-based classifier called simplified fuzzy ARTMAP (SFAM). Figure 33 illustrates the structure and the learning scheme of SFAM. The neural network learns about the “normal” and the “abnormal” traffic patterns from a set of training data, which are the feature vectors generated by the sensors. In Figure 33, these training vectors are depicted as the green (“normal”) and the red (“abnormal”) points in the multidimensional feature space. When the features are chosen properly, points with the same color tend to form clusters. The SFAM identifies the clusters and builds hypercubes around them. Each hypercube is associated with either the “normal” or the “abnormal” output class. Given a new feature vector, the classifier evaluates the proximity of the vector to the hypercubes and assigns it to the class that the closest cube belongs to.




Figure 33. The structure and the learning scheme of simplified fuzzy ARTMAP classifier.

5.4 USER INTERFACE

A graphical user interface, shown in Figure 34, allows a user to change the settings of the vehicle sources and control the learning process. Through the interface, the user can also observe the traffic flow, monitor the system output, and examine the internal parameter values such as the feature vectors and the weights of the neural network.

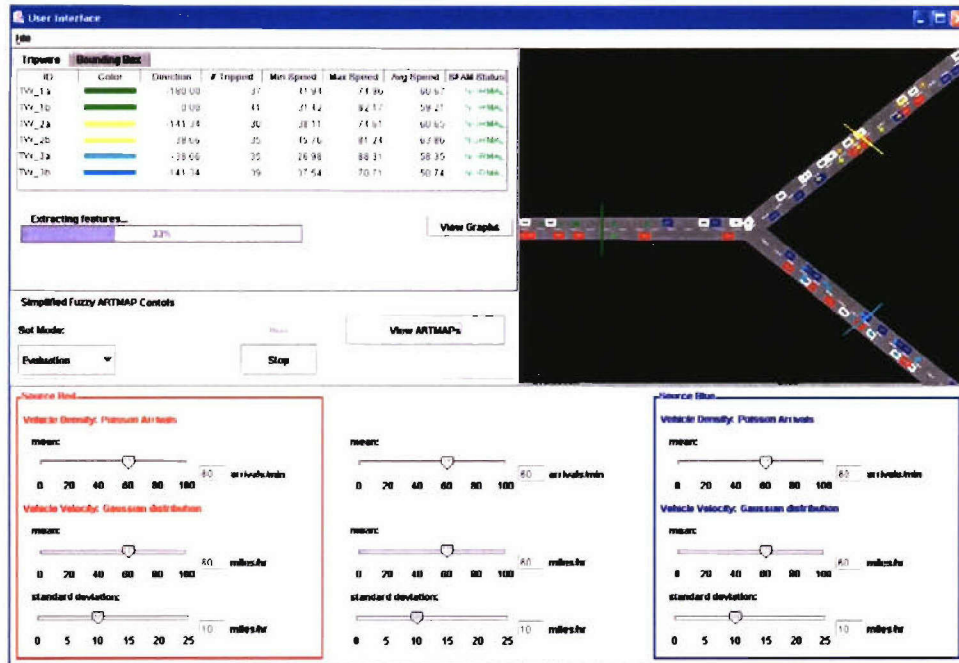


Figure 34. Test bed graphical user interface. The three bottom panels control the red, white, and blue vehicle sources. The upper-right panel displays the traffic flow and sensor locations (tripwires shown). The upper-left panel shows the internal parameters such as the feature vector values and the status of the traffic (normal or abnormal). The middle-left panel is the control for SFAM.

5.5 SOFTWARE ARCHITECTURE

The software design of the test bed is shown in Figure 35. It has a “publish-subscribe” software architecture. A centralized communication manager supports broadcast communications.

Since data transfer among different parts of the system is mediated by the communication manager, it is easy to swap components in and out of the system. For instance, if a “convoy detector” needs to be added as a new feature or some new detectors become available, they can be easily plugged into the system. These new additions are illustrated in red in Figure 35.

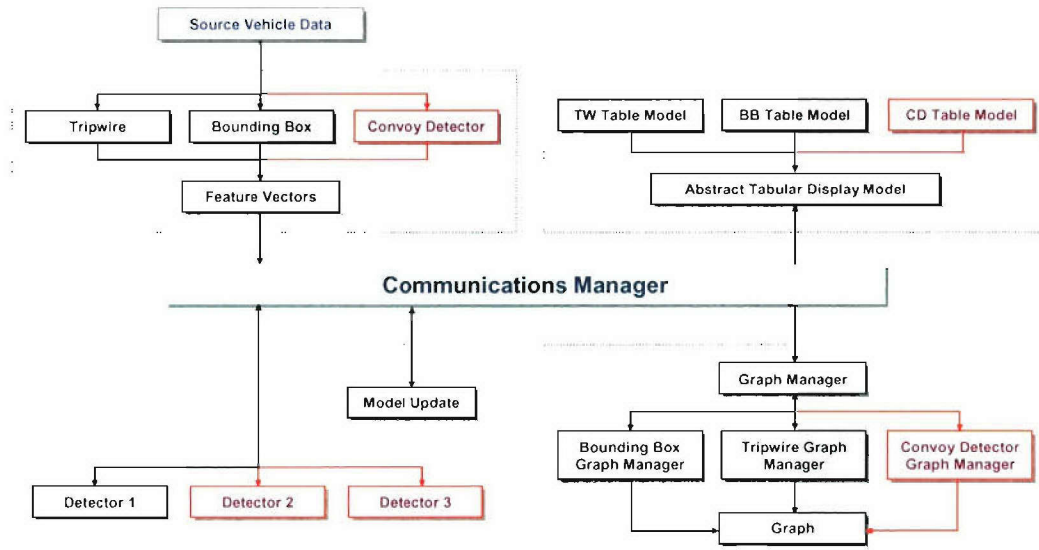


Figure 35. Test bed “publish-subscribe” software design. The upper-left part hosts the sensors. The upper-right controls the tabular displays of parameters. The lower-left contains the detectors. The lower-right provides all the graphical displays.

6. SUMMARY

Traditional target tracking may not be the best method to exploit GMTI data for large area persistent surveillance. In this work, we consider the approach that uses the GMTI data as moving spots on the ground to estimate the level of activities in an area.

A computational framework is proposed for data processing and inference. This framework has a bottom-up, hierarchical, and modular structure where sensor data provide input to the system at the bottom level. The system design emphasizes on evidence accumulation and continuous learning. Various pattern recognition, machine learning, and data mining algorithms can be implemented in the processing modules of the system. The modules at different levels share a common functional structure: after preprocessing, computational features are extracted for model building and reasoning; the results are then passed to the modules at the next level up. This computational design ensures that the system is easily extendable and can be tested using a generalized test bed.

Traffic data from the ADMS Virginia database were used as a surrogate of GMTI data. We explain in detail how low-level features are constructed from the traffic data and used as indicators to the activity level at Norfolk Naval Base. A convoy detection algorithm for exploiting GMTI data is also described.

A test bed was built for evaluating detection and reasoning algorithms. It allows easy change of input, output, and algorithm plug-ins as well as access to internal parameters.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</small>					
1. REPORT DATE (DD-MM-YYYY) 02-10-2004		2. REPORT TYPE Technical Report		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE Activity Level Change Detection for Persistent Surveillance				5a. CONTRACT NUMBER FA8721-05-C-0002	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Fang Liu and Lawrence Bush				5d. PROJECT NUMBER 1204	
				5e. TASK NUMBER 506	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MIT Lincoln Laboratory 244 Wood Street Lexington, MA 02420-9108				8. PERFORMING ORGANIZATION REPORT NUMBER TR-1104	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AF Space and Missile Systems Center/YS 2420 Vela Way, Suite 1467-A8 El Segundo, CA 90245				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) ESC-TR-2005-061	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT A new approach to GMTI data exploitation for large area persistent surveillance is presented. Instead of traditional target tracking, this approach utilizes GMTI data as moving spots on the ground to estimate the level of activities and detect unusual activities such as military deployments. A multilayer hierarchical exploitation scheme is proposed. This computational framework has clean interfaces between layers consisting of multiple processing modules. Various data processing, machine learning, and reasoning algorithms can be implemented in these modules. This system is easily extendable and can be tested using a generalized test bed. The development of two processing modules, vehicular volume and convoy detector, is described. For the vehicular volume module, US highway data were used as a surrogate of long-term GMTI surveillance data. The relationship between the activity level of Norfolk Naval Base and the traffic pattern on a road leading to the Base is studied. The convoy detection module, developed using real GMTI data, contains an algorithm that detects convoys without explicit target tracking. An end-to-end testing facility was also developed. Using this test bed, the system can be tested at different levels: as an individual processing module, as multiple cooperating processing modules across layers, or as the entire system.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code)
			None	50	